

Estimating Activity at Multiple Scales using Spatial Abstractions

Majd Hawasly¹, Florian T. Pokorny², and Subramanian Ramamoorthy¹

¹ School of Informatics, University of Edinburgh

² KTH Royal Institute of Technology

Abstract. Autonomous robots operating in dynamic environments must maintain beliefs over a hypothesis space that is rich enough to represent the activities of interest at different scales. This is important both in order to accommodate the availability of evidence at varying degrees of coarseness, such as when interpreting and assimilating natural instructions, but also in order to make subsequent reactive planning more efficient. We present an algorithm that combines a topology-based trajectory clustering procedure that generates hierarchically-structured spatial abstractions with a bank of particle filters at each of these abstraction levels so as to produce probability estimates over an agent’s navigation activity that is kept consistent across the hierarchy. We study the performance of the proposed method using a synthetic trajectory dataset in 2D, as well as a dataset taken from AIS-based tracking of ships in an extended harbour area. We show that, in comparison to a baseline which is a particle filter that estimates activity without exploiting such structure, our method achieves a better normalised error in predicting the trajectory as well as better time to convergence to a true class when compared against ground truth.

1 Introduction

Autonomous agents that act in dynamic environments, including humans and other active agents, must possess the capacity to make predictions about the changes in the environment and correspondingly take actions with respect to their best estimate about the state of the world. In particular, there is an increasing need for robots to make predictions about the activity of other agents - what is the state of the activity of the other agent and what might they do next?

Traditionally, these prediction problems have been addressed using tools from *state estimation theory*, which is most mature when the source of uncertainty is rooted in the dynamics and noise of sensorimotor processes. Particle filters are used extensively in robotics when there is the need to perform Bayesian state estimation in nonlinear systems with noisy and partial information about the underlying state [1,2,3]. In this approach, the posterior probability distribution over the states given a sequence of measurements is approximated with a set of particles which represent state hypotheses.

When the hypotheses pertain to spatial activity, such as navigation of people/robots in typical human-centred environments, or perhaps in more complex configuration spaces such as when dealing with full body motion, the underlying dynamics are best described in a hierarchical fashion, so that a movement is not just determined by the local laws of dynamics and noise characteristics, but also by longer term goals and preferences. This has at least two important implications for predictive models. We need techniques such as for activity estimation to be able to accept evidence at varying scales - ranging from very precise position measurements to coarser forms of human feedback (e.g., “she is heading to the right between two obstacles”) [4,5] or variable resolution sensory signals (e.g., GPS-like devices that provide location estimates to within a spatial neighbourhood of variable extent). Correspondingly, we would like to be able to output predictions at multiple scales, to support decision making at all these levels. These desiderata form the primary focus of this paper, in which we propose a technique that accepts evidence and provides estimates at multiple resolutions.

There is a long tradition of hierarchical modelling of motion which could inform the design of such techniques. Early models of large-scale spatial navigation [6] considered ways in which multiple representations, ranging from coarse and intuitive topological notions of connectivity between landmarks to a more detailed metrical and control level description of action selection, could be brought together in a coherent framework and implemented on robots. Other recent methods, driven more from motion planning and control considerations e.g. [7,8], propose ways in which control vector fields could be abstracted so as to support reasoning about the hybrid system that is aimed at solving the larger-scale tasks.

While these works provide useful inspiration, they do not directly address our aforementioned desiderata. Firstly, the hierarchy is often statically defined by the designer of the system and the algorithm. In many applications, it is of interest to be able to learn these from data - both because this enables online and continual adaptation over time, but also because the underlying principles determining the types of motion may not be fully understood (e.g., the activity of people in a complex environment, driven by private and varied utility functions). Secondly, the approaches that are principled in the way they define the hierarchy are often silent on how best to integrate with the methodology for maintaining Bayesian belief estimates, such as with a particle filter - the question of how best to define a correspondingly-hierarchical activity estimation method is largely open.

There is indeed prior work on the notion of hierarchy in state estimation with particle filters. For instance, Verma et al. [9] define a variable resolution particle filter for operation in large state spaces, where chosen states are lumped into aggregated states so that the complexity of the particle filter may be reduced. Brandao et al. [10] devise a subspace hierarchical particle filter wherein the focus is on defining subspaces for which state estimation calculations can be run in parallel, alleviating the computational burden through factored parallel computation. There are a variety of ways in which computations have been factored [11], e.g., by partitioning the ways in which sampling calculations may

be performed for tracking articulated objects with an implied hierarchy [12], of by using a hierarchy of feature encodings [13] and contour detection [14].

In this paper, we take a different approach to learning the hierarchy, using which we devise a novel construction of a bank of particle filters - one at each scale - maintaining consistent beliefs over the trajectories as a whole (in the spirit of plan and activity recognition) and through that over the state space. Given a set of trajectories (such as from historical observations of activity) and a notion of trajectory similarity, we define procedures for hierarchical clustering of these trajectories. The output of clustering is a tree-structured representation of trajectory classes that correspond to incrementally-coarser partitions of the underlying space. We present an agglomerative clustering scheme [15] using the Fréchet distance between trajectories [16]. This construction of trajectory clustering in the form of a filtration is inspired by earlier work using persistent homology [17,18,19], but is instantiated in a simpler and computationally more efficient manner through the use of Fréchet distance based agglomeration. Equipped with this data-driven notion of a hierarchy, we show how to define the dynamics at the different levels, and how they can be employed with a new stream of observations to provide probability updates over time and over the classes in the tree.

Our construction of the filter allows us to fluently incorporate readings of varying resolution if they were accompanied by an indication of the coarseness with which the observation is to be interpreted. This issue of variability is much broader in scope, covering many other aspects of dynamical systems behaviour [20,21].

We evaluate our proposed method by showing how unsupervised learning of hierarchical structure in the activity data enables the bank of particle filters at multiple scales (which we refer to more concisely, with slight abuse of terminology, as a ‘multiscale particle filter’) to perform better than baselines both in terms of normalised error in predicting the position of an agent with respect to the ground truth trajectory, and in terms of the time taken for the belief to converge to the true trajectory of class (depending on the resolution of the prediction being considered). We perform such experiments first with a synthetic dataset which brings out the qualitative behaviour of the procedure in a visually intuitive manner, and then with a real world dataset based on tracks of ships in a harbour (based on a database associated with the worldwide AIS system).

2 Hierarchical abstraction of trajectories

To create the filtration of spatial abstractions from trajectories, we consider hierarchical clustering [22] by means of discrete Fréchet distance [16]. For two discretised d -dimensional trajectories $\tau_1 : [0, m] \rightarrow \mathbb{R}^d$ and $\tau_2 : [0, n] \rightarrow \mathbb{R}^d$:

$$\delta_F(\tau_1, \tau_2) = \inf_{\alpha, \beta} \max_{j \leq m+n} \delta_E(\tau_1(\alpha(j)), \tau_2(\beta(j))), \quad (1)$$

where α and β are discrete, monotonic re-parametrisations $\alpha : [1 : m+n] \rightarrow [0 : m]$, $\beta : [1 : m+n] \rightarrow [0 : n]$ which align the trajectories to each other

point-wise, and $\delta_E(.,.)$ is the Euclidean distance between two points. This metric δ_F corresponds to the maximal point-wise distance between two optimal reparameterisations of τ_1, τ_2 , and it can be computed efficiently using dynamic programming in $O(mn)$ time [16].

After computing the distance matrix D for the trajectories, where $D_{i,j} = D_{j,i} = \delta_F(\tau_i, \tau_j)$ and $D_{i,i} = 0$, the trajectories can be considered as data points to be clustered, $C = \{\tau_1, \tau_2, \dots, \tau_M\}$. A hierarchical agglomerative clustering [15] of C results in a tree of trajectory clusters (see Figure 2).

Hierarchical agglomerative clustering is an iterative approach to data clustering in which, at every iteration, two clusters at a lower level get merged to make a single new cluster. This gives rise to a tree data structure in which the leaves at one end are the individual data items, and the root node at the other end is the cluster made by merging all data points together, while the intermediate layers combine data items based on their proximity. The order of merging depends on the distance between the clusters, such that the pair with the smallest distance are merged first. Thus, every new cluster can be assigned a distance value at which it gets created. An important consideration in the design of hierarchical agglomerative clustering algorithms is the method of computing the distance between clusters.

A *single linkage* algorithm assumes the distance between two clusters to be the smallest distance between the individual data items in the two clusters. At stage t , for a collection of objects C^t and a distance matrix D^t , the pair of distinct clusters with the smallest distance in D^t are merged to create a new cluster, $\tau_{ij} = \cup_h \tau_h$ for $h \in \arg \min_{i,j,i \neq j} D_{i,j}^t$. C^{t+1} is created by removing the two clusters τ_i and τ_j and adding τ_{ij} . The distance matrix is also updated to reflect the change, $D_{ij,h}^{t+1} = \min(D_{i,h}^t, D_{j,h}^t)$ for all $\tau_h \in C^t$. The process repeats until it terminates when only one cluster remains.

We call $D_{i,j}$ the *birth index* of cluster τ_{ij} , and we denote it by b_{ij} . This refers to the distance threshold after which τ_{ij} starts to exist. Also, we define the *death index* $d_i = d_j = D_{i,j}$ to be the distance index at which cluster τ_i (similarly, τ_j) ceases to exist.

The output of this algorithm is a tree structure $\mathcal{T}(\mathcal{C}, \rho)$ where $\mathcal{C} = \bigcup_t C^t$ is a collection of all the hierarchical clusters (original ones included), and the parent function $\rho : \mathcal{C} \rightarrow \mathcal{C}$ which maps a cluster to its immediate *parent*. Then, if $\rho(\tau_i) = \tau_h$, then $\tau_i \subset \tau_h$ and $b_h = d_i$; see the tree in Figure 2.

We consider a tree node $\tau_i \in \mathcal{C}$ to be *alive* at some birth index b if $b_i \leq b < d_i$, i.e., when it is born but not yet dead. A *level* in the tree, identified with a birth value b , contains all the nodes that are alive at b . We denote the level at index b by $\mathcal{C}_b \subseteq \mathcal{C}$.

Figure 1 illustrates clustering 14 representative trajectories of navigation around a Y-junction using the method described above. Each panel shows the newly-created cluster in one level of the hierarchy at some birth index. The rest of the trajectories appear in grey in the background.

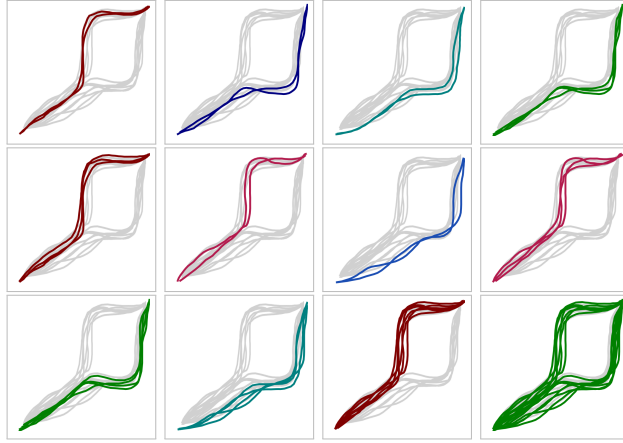


Fig. 1: Clustering a set of 14 trajectories using hierarchical single-linkage agglomerative clustering with Fréchet distance. Panels show the new clusters at different levels of the filtration (tree levels). Note that the rest of the trajectories/clusters at each level are depicted in the background for the sake of clarity.

3 Multiscale Hierarchy of Particle Filters

A Bayesian particle filter [23] tracks a probability distribution (a *belief*) over some random variable of interest $z^t \in \mathbb{R}^d$ by evolving a collection of N hypotheses called *particles* utilising prior knowledge $\mathbf{P}(z^0)$ and a sequence of measurements $\xi^{1:t}$.

Upon receiving a new observation ξ^t , a Bayesian belief should be updated as follows:

$$\mathbf{P}(z^t | \xi^{1:t}) \propto \mathbf{P}(\xi^t | z^t) \sum_{z^{t-1} \in \mathbb{R}^d} \mathbf{P}(z^t | z^{t-1}) \mathbf{P}(z^{t-1} | \xi^{1:t-1}).$$

Sampling directly from the target distribution might not be feasible, so a particle filter computes an approximation that involves representing its beliefs by a set of particles. In the standard algorithm [3], particles are sampled from a *proposal distribution* (typically, the dynamics model $\mathbf{P}(z^t | z^{t-1})$) and the deficit between the two distributions is rectified by assigning *importance weights* $w^t \in [0, 1]$ to the particles (typically, the observation likelihood $\mathbf{P}(\xi^t | z^t)$). The actual belief is retained again by *resampling* N particles according to their weights. In practice, one replaces a small fraction of all particles randomly with new ones regardless of their weights in a bid to avoid particle *depletion*.

We present the Multiscale Hierarchy of Particle Filters (MHPF), a stack of *consistent* particle filters defined over *abstractions* of the value of the random variable. In this paper, the random variable of interest is the agent’s navigation plan, encoded quantitatively in its point position, and qualitatively in the class and shape of its planned trajectory. We assume that a collection of trajectories are available to MHPF in order to construct this abstraction from data. The

abstractions are representations of this random variable at decreasing resolution, so that the lowest level of the abstractions hierarchy consists of complete trajectories at the smallest scale (with cardinality equal to the size of the trajectory dataset). At any higher level, these trajectories are clustered into a smaller number of bins or categories, representing coarser descriptions of the trajectory shapes. Thus, at any stage the status of the process can be queried at any of the different levels of resolution. This gives the additional advantage of allowing evidence at various degrees of coarseness to be incorporated into the filter by using it to update the probability estimate at that level. In order to maintain consistency between the particle filters of the stack, this update results in corresponding updates to all the other filters - based on a procedure to be described below.

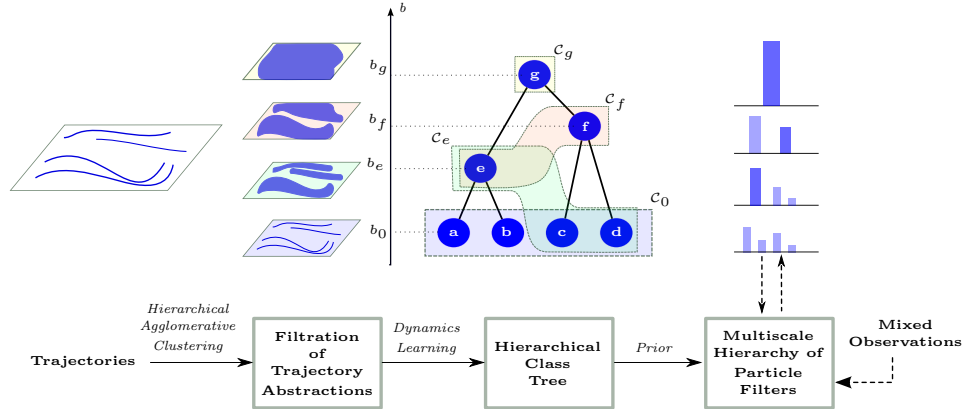


Fig. 2: An overview of the approach. A collection of trajectories are clustered at different resolutions to give a filtration of spatial abstractions, and dynamics models are acquired for the individual classes. The classes are organised in a tree structure with their birth indices. The shaded areas show levels of the hierarchy, where C_0 is the finest level with single trajectory classes, and C_g is the coarsest class, representing all the trajectory data combined. The distributions on the right show an example of a consistent distribution across the tree. A collection of particle filters encode these distributions and allow coarse observations to be handled at any level of the hierarchy.

3.1 Setup

Given a tree $\mathcal{T}\langle\mathcal{C}, \rho\rangle$ over the input trajectories, each cluster $c \in \mathcal{C}$ is a collection of ‘similar’ trajectories at some level of resolution b . The cluster c then can be seen as a *class* of behaviour for the tracked process, identified by a generative dynamics model $\theta = \mathbf{P}(z'|z, c)$, $z, z' \in \mathbb{R}^d$ from which the member trajectories are samples.

The dynamics of a class c are approximated from the points of the member trajectories using a localised model as follows. All the points of c in an ϵ -ball

around the point of interest $z \in \mathbb{R}^d$ are located $B_\epsilon(z) = \{z' \in \mathcal{C} : \delta_E(z', z) < \epsilon\}$, and the local velocities at these points \dot{z}' are used to estimate the new velocity, $\dot{z} = \eta \sum_{z' \in B_\epsilon(z)} 1/\delta_E(z, z') \dot{z}'$, where $\eta^{-1} = \sum_{z' \in B_\epsilon(z)} 1/\delta_E(z, z')$ is a normalisation factor. Then, a new position is sampled, $z^* \sim z + \dot{z} + \gamma(\kappa)$, where $\gamma(\kappa)$ is a noise term related to dynamics noise parameter κ .

The tree of clusters \mathcal{T} and the associated dynamics are the input to MHPF. As a stack of filters, a distinct filter is defined for every level of the tree \mathcal{T} . Thus, a particle x^t in MHPF is a weighted hypothesis of the *class and the position* at time t . That is, every particle at some level b represents a hypothesis not only for the position of the tracked process in \mathbb{R}^d but also which of the different classes in \mathcal{C}_b represents the behaviour of the process best. We write $(x^t(z^t, c), w^t)$ where $z^t \in \mathbb{R}^d$ is the position, $c \in \mathcal{C}_b$ is the class, and w^t is a weight that reflects to what extent the hypothesis of the particle is compatible with the evidence. We denote by X_b the set of particles at level b .

There are two kinds of observations in MHPF: 1) the noisy position observations $z^t + \gamma(\psi)$, where $\gamma(\psi)$ is a noise term related to the observation noise parameter ψ , which are the typical observations for standard particle filters as well as the filter of the lowest level of the MHPF stack; and 2) *coarse observations* which provide evidence regarding the underlying process and can be identified to one of the classes in \mathcal{C} other than \mathcal{C}_0 ¹. In both cases, the MHPF returns a stack of consistent probability distributions pertaining to the different tree levels.

3.2 Procedure

MHPF is based on a probability distribution defined at the finest level \mathcal{C}_0 from which the tree is rebuilt, as shown in the procedure in Algorithm 1.

First, the particle set X_0 is created by sampling N particles from a prior $\Delta(\mathcal{C}_0 \times \mathbb{R}^d)$ over class assignment and initial positions, then assigning them equal weights, where $\mathcal{C}_0 = \{c \in \mathcal{C} \mid b_c = b_0\}$ is the collection of individual trajectories forming a class each at the lowest clustering threshold b_0 . Denote by N_i the number of particles from class c_i , such that $\sum_{c_i \in \mathcal{C}_0} N_i = N$.

The probabilities of the classes of \mathcal{C}_0 are computed from the initial weights, and these probabilities in turn are used to compute the probabilities of the rest of the classes as described in Algorithm 2.

At this stage, the class probabilities are propagated recursively upwards by the additivity rule, so that a parent's probability is the sum of its children's probabilities, $\mathbf{P}^t(\bar{c}) = \sum_{\underline{c} \in \rho^{-1}(\bar{c})} \mathbf{P}^t(\underline{c})$. In order to understand the intuition behind this step, consider the probabilities assigned to the classes/nodes of the tree with respect to the *regions* that are defined by a spatial nearest neighbour relationship to the points of their corresponding trajectories. Consider the example of a 2-dimensional domain in Figure 3, where the region corresponding to a class can be understood as the union of 2-dimensional Voronoi cells of a discretisation


¹ This is compatible both with variable resolution sensors (e.g. GPS receivers) and with high-level qualitative instructions (e.g. linguistic instructions) as long as a mapping can be established between the observation and \mathcal{C} , especially in the latter type.

Algorithm 1 Multiscale Hierarchy of Particle Filters


Require: Prior over particles, number of basic particles N , the depletion parameter v , tree structure \mathcal{T}

- 1: Sample N particles from the prior $\Delta(\mathcal{C}_0 \times \mathbb{R}^d)$ with equal weights.
 - 2: **for** each time step $t > 0$ **do**
 - 3: Build the tree probabilities up from \mathcal{C}_0 (Algorithm 2).
 - 4: **for** parents \bar{c} of \mathcal{C}_0 recursively **to** the top of \mathcal{T} **do**
 - 5: Sample a number of particles from \bar{c} ; $N_{\bar{c}} = \sum_{c=\rho^{-1}(\bar{c})} N_c$, with equal weights.
 - 6: **end for**
 - 7: **for** each particle $x(z^{t-1}, c)$ **do**
 - 8: sample a new position $z^t \sim \mathbf{P}(z|z^{t-1}, c)$
 - 9: **end for**
 - 10: Receive observation ξ^t .
 - 11: Update tree with observation ξ^t (Algorithm 3).
 - 12: Resample $N(1 - v)$ particles from \mathcal{C}_0 based on the updated weights w^t .
 - 13: Sample Nv particles uniformly randomly from \mathcal{C}_0 .
 - 14: **end for**
-

of the class trajectories. Thus, merging two classes in the tree is analogous to merging the regions associated with their classes, and correspondingly adding the probability of the two child classes to yield the probability of the parent. Similarly, the children of a class proportionally inherit their parent's probability when moving downward in the tree. The Voronoi cells depicted below are never explicitly computed. However, implicitly, this defines our notion of consistency between the probability estimates at the levels of the hierarchy.



(a) Three separate classes



(b) Two classes merge

Fig. 3: An illustration of the intuition behind the relationship between the spatial abstraction and the corresponding probability operations on the tree in a toy example 2-dimensional domain.

With the probabilities specified, the same number of particles as the total number of children's particles are created for the parent \bar{c} , $N_{\bar{c}} = \sum_{c=\rho^{-1}(\bar{c})} N_c$, and this is repeated recursively to the top of the tree. Note that any arbitrary level b of the tree would have exactly N particles with a proper probability distribution, while the total number of particles in the full tree depends on how the particles are distributed between the tree branches. The last stage of the

tree construction is to sample new positions for the particles. Note that the class label of an individual particle does not change by sampling.

Algorithm 2 Tree Probability Rebuild

Require: Tree structure \mathcal{T} , tree level \mathcal{C}_b , particle set X_b

1: Update the probabilities of classes $c_i \in \mathcal{C}_b$ from X_b weights,

$$\mathbf{P}^t(c_i) = \frac{\sum_{c(x)=c_i} w^t(x)}{\sum_{x \in X_b} w^t(x)}$$

2: **for** children of \mathcal{C}_b recursively **to** the bottom of \mathcal{T} **do**

3: Update the probability of the child \underline{c} relative to the probability of its parent $\bar{c} = \rho(\underline{c})$:

$$\mathbf{P}^t(\underline{c}) = \mathbf{P}^{t-1}(\underline{c}) \frac{\mathbf{P}^t(\bar{c})}{\mathbf{P}^{t-1}(\bar{c})}$$

4: **end for**

5: **for** parents of \mathcal{C}_0 recursively **to** the top of \mathcal{T} **do**

6: Update the probability of the parent \bar{c} relative to the probability of its children $\underline{c} = \rho^{-1}(\bar{c})$:

$$\mathbf{P}^t(\bar{c}) = \sum_{\underline{c}=\rho^{-1}(\bar{c})} \mathbf{P}^t(\underline{c})$$

7: **end for**

3.3 Observations

A coarse observation $\xi^t \in \mathcal{C}$ of level b_ξ relates to all the particles from classes that are *alive* at that level of the hierarchy, $\mathcal{C}_\xi = \{c_i \in \mathcal{C} \mid b_i \leq b_\xi < d_i\}$.




Fig. 4: A coarse observation targets a layer in the tree, and the particles of all the classes that are alive will be updated using tree class distance.

To update a particle $x(z^t, c)$ of a class $c \in \mathcal{C}_\xi$ we use the *tree class distance* between c and ξ^t , which we define for two classes c_1 and c_2 as the birth index of the first shared parent of c_1 and c_2 in the tree. This distance measures how far we have to climb in the tree for the two classes to be similar enough and join the same cluster, or alternatively how large the ϵ -balls around the points of one class need to be to include the other. For example, the class distance between class e

and class g in Figure 2 is b_g , which is also the case for classes c and e . The weight of a particle is then updated relative to the tree distance, $w \propto -\log(\delta_{\mathcal{T}}(\xi^t, c))$.

On the other hand, updating a particle with a position observation is straightforward, relative to the Euclidean distance between the observation and the particle's position, $w \propto -\log(\delta_E(\xi^t, z))$.

Algorithm 3 Observation Update

Require: observation ξ^t , tree structure \mathcal{T}

```

1: if fine observation then
2:    $\hat{\mathcal{C}} = \mathcal{C}_0, \hat{X} = X_0$ 
3:   for every particle  $x(z, c) \in \hat{X}$  do
4:     update weight relative to Euclidean distance to  $\xi^t$ :  $w^t \propto -\log(\delta_E(z, \xi^t))$ .
5:   end for
6: else if coarse observation then
7:   Identify the level  $b_\xi$  of the observation  $\xi^t$ 
8:   Find all the classes that are alive at the level:  $\hat{\mathcal{C}} = \{c_i \in \mathcal{C} : b_i \leq b_\xi < d_i\}$ .
9:   for every class  $c \in \hat{\mathcal{C}}$  do
10:    find the tree distance  $\delta_{\mathcal{T}}(c, \xi^t)$ , the birth index of the first shared parent of  $c$ 
    and  $\xi^t$ .
11:    for every particle  $x(z, c) \in X_c$  do
12:      update weight relative to tree distance to  $\xi^t$ :  $w^t \propto -\log(\delta_{\mathcal{T}}(c, \xi^t))$ .
13:    end for
14:  end for
15: end if
16: Rebuild the tree probabilities from  $\hat{\mathcal{C}}$  (Algorithm 2).
17: for every particle  $x(z, c) \in X_{\hat{\mathcal{C}}}$  do
18:   Update particle weights:  $w^t = w^{t-1} \frac{\mathbf{P}^t(c)}{\mathbf{P}^{t-1}(c)}$ 
19: end for

```

Then, the probabilities of the classes of \mathcal{C}_ξ are recomputed as the sum of their particles' normalised weights. Note that the coarse update is qualitative in nature such that all the particles of a certain class from \mathcal{C}_ξ would get the same update regardless of the particle positions.

The updated class probabilities of \mathcal{C}_ξ propagate to the rest of the tree as in Algorithm 2. At this stage, children of updated classes are updated first recursively relative to their parents' new probabilities,

$$\mathbf{P}^t(\underline{c}) = \mathbf{P}^{t-1}(\underline{c}) \frac{\mathbf{P}^t(\bar{c})}{\mathbf{P}^{t-1}(\bar{c})}, \quad \forall \underline{c} = \rho^{-1}(\bar{c}).$$

Then, the updates propagate upwards by updating all parents recursively, summing up their children's probabilities,

$$\mathbf{P}^t(\bar{c}) = \sum_{\underline{c}=\rho^{-1}(\bar{c})} \mathbf{P}^t(\underline{c})$$

Once the tree probabilities are balanced, the rest of the particle weights are updated to reflect their new class probabilities, $w^t = w^{t-1} \frac{\mathbf{P}^t(c)}{\mathbf{P}^{t-1}(c)}$, $\forall x(., c) : c \in \mathcal{C} \setminus \mathcal{C}_\xi$.

The final step is to resample N particles from the finest level of the tree X_0 with equal weights to get the posterior particle set after incorporating the evidence. To guard against particle depletion, we randomly replace the classes of v of the particles uniformly randomly to classes from \mathcal{C}_0 . From this new particle set the process repeats.

4 Experiments

We evaluate the performance of **MHPF** in a number of 2-dimensional navigation domains over two baselines which are particle filters without access to the hierarchical structure. **BL1** is a basic particle filter [3] with N particles $x(z, c)$, with c restricted to the single trajectory class \mathcal{C}_0 . Thus, each particle follows its single-trajectory class. Secondly, **BL2** is a particle filter with N particles $x(z)$ where all the particles follow the localised dynamics of the combination of all the trajectories with κ noise. Note that **BL1** is equivalent to the bottom layer in the MHPF filter stack, and **BL2** is equivalent to the top layer.

We show the improvements using a number of metrics. We use the mean squared error of the filter’s point prediction to show efficacy, and we evaluate performance by showing the distance of the filter’s predicted class to the ground truth as well as the time needed to converge to the true class being followed by the agent.

We use synthetic datasets as well as real world data in the experiments. Each experiment runs over 10 randomly selected scenarios described by corresponding ground truth trajectories that the process follows. The trajectories are uniformly discretised, and the length of a trial will depend on the number of points in the discretisation. Each trial is repeated 25 times and the results are averaged.

Each of the experiments had $N = 100$ particles at any of the tree levels. At every time step, observations are generated from the discretised ground truth. A fine observation is defined as $z^t + (\gamma_1, \gamma_2)$ where $z^t \in \mathbb{R}^2$ is ground truth at time t , $\gamma_1, \gamma_2 \in \mathbb{R}$ are chosen uniformly randomly from $[0, \psi]$ where ψ is the observation noise parameter of the experiment. A coarse observation is generated by sampling a number of points ($n = 10$) from $\mathcal{N}(z^t, \psi^2)$ where \mathcal{N} is a normal distribution, z^t is the ground truth at time t and ψ is the observation noise parameter. Then, the class that has the highest probability to generate these samples is chosen as ξ^t . For dynamics we used a localised model as in Section 3.1 with ϵ (the size of the ϵ -ball) equals to b_c for some coarse class c and with a noise term from $[0, \kappa]$ for the dynamics noise parameter κ . We used KD-trees for efficient selection of neighbourhood points. At the end of every step, $v = 1\%$ of the particles is changed randomly.

4.1 Synthetic datasets

We work with two synthetic domains, the first represents a 2-dimensional configuration space with 33 trajectories with general start and end positions, and the second one has 13 trajectories with a fixed start and end position (Figure 5).




Fig. 5: Synthetic datasets of 33 trajectories in a 2D configuration space with obstacles, and 13 trajectories that starts and ends at the same position in a 2D plane.

For the configuration space dataset, we compute the filter’s predicted position at time t as the w -weighted average of the particle positions, and report the average of the mean squared error (MSE) of the ground truth to this predicted position over time and for 10 random scenarios, each repeated 25 times. **MHPF** achieved a mean of 0.27 (standard deviation of 0.04) compared to **BL1** which achieved 0.38(0.14) and **BL2** which achieved 0.53(0.13). This experiment uses fine observations only.

Figure 6 illustrates the kind of multi-resolution output the filter can produce. It shows the evolution of the filter’s *maximum a posteriori* (MAP) class with time and for different levels of the tree. Each column shows the classes of some level b , with the leftmost column showing the finest level with individual trajectory classes and the rightmost column showing the coarsest level (a single class combining all the trajectories), while rows show progress over time. The thicker the trajectories are, the more likely their class is.

Next, using the 13 trajectory dataset, we compare **MHPF** with **BL1** in a situation where, in addition to the consistent fine observations, coarse observations are produced stochastically 50% of the time. This is motivated by use cases where high-level qualitative information (e.g. human instructions) might exist along the finer localised measurements. We analyse the benefit of this additional knowledge by plotting the average tree class distance of the MAP prediction of the filters to the ground truth. We show the results for different values of dynamics noise ($\kappa = 30\%, 50\%, 75\%$) and different values of observation noise ($\psi = 1\%, 5\%$). The results are reported in Figure 7a.

Finally, using the same dataset, we analyse the situation where fine observations are only provided for a fraction of the time (a lead-in period of 5% and 7.5% of the trial length), then only coarse observations are given. We present the effect of that on the time needed by **MHPF** and **BL1** to converge by plotting the time needed for the class distance to reach within the 33%-ball of the ground

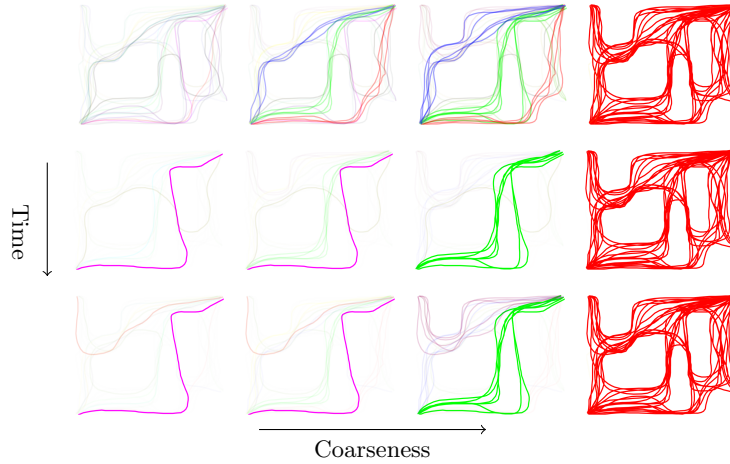


Fig. 6: Evolution of the filter’s prediction. The columns are for different levels across the tree with the finest at the left. The rows are for different time steps with the first at the top. Each panel shows the trajectories of the alive classes. The opacity of the line reflects the probability of the class.

truth. We show the results for observation noise $\psi = 1\%$ and for different values of dynamics noise ($\kappa = 30\%, 75\%$). The results are reported in Figure 7b.

4.2 Tanker vessel data

This experiment uses publicly-available data regarding the movement of ships in a harbour area. Specifically, we utilise records of tanker vessel tracks around the Gulf of Mexico [24]. From the data which is available in the form of density/occupancy graph we generate 194 trajectories by weighted random walks from manually-selected initial positions, such that a trajectory is more likely to follow the denser areas and does not change direction much often. Figure 8 shows the density and the trajectory classes for $b = 70$.

Compared to **BL1**, we explore the benefit to convergence when receiving coarse information 50% of the time along with fine observations in ship tracking scenarios. The reported values in Figure 9 are averages of tree class distance between the MAP prediction of the filter and the ground truth trajectory with dynamics noise ranging from 10% to 30% and observation noise ranging from 10% to 20%.

5 Conclusion

We propose a novel approach to utilising a hierarchical clustering over trajectories (a filtration) to devise a correspondingly hierarchical representation of probability distributions over the underlying state space so as to enable Bayesian

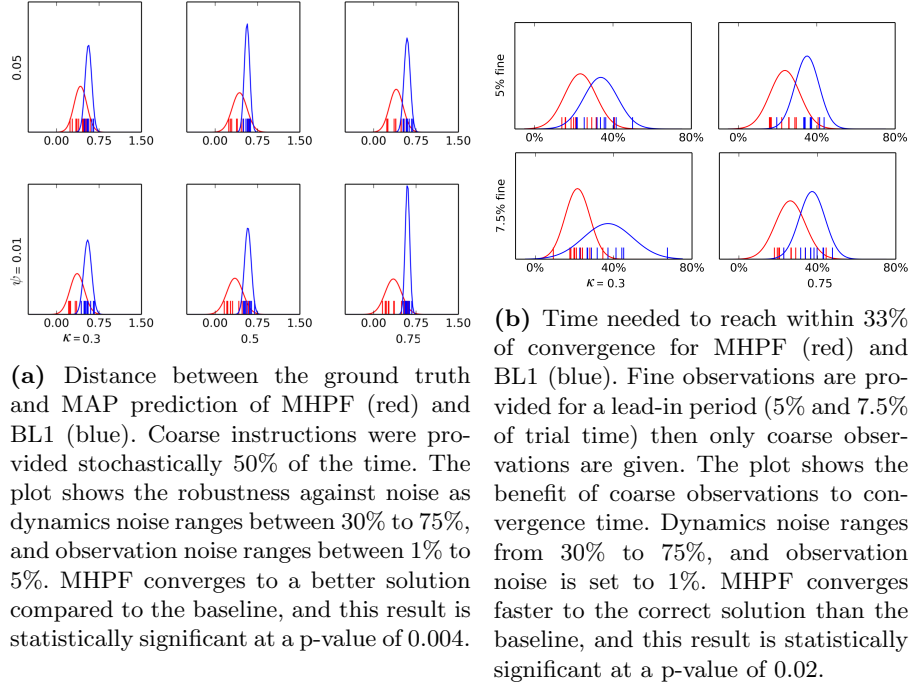


Fig. 7: Performance results in the synthetic domains comparing MHPF (red) and BL1 (blue). X axis shows dynamics noise. The experiments were run on 10 different test scenarios, each repeated 25 times, then the results were averaged.




Fig. 8: Density and the trajectory classes at one level in the tree in the tanker vessel experiment

filtering. A key benefit of our methodology is the ability to incorporate coarse observations in the estimation process to seamlessly allow for potential inhomogeneity in sensor readings, such as when a GPS device obtains position fixes with varying confidence, or for signals at varying degrees of coarseness, such as when a human user instructs a robot in relational terms. We demonstrate the usefulness of this technique with experimental domains of increasing complexity, ranging from a synthetic data set intended to illustrate the elements of the




Fig. 9: Distance between the ground truth and the MAP prediction of MHPF (red) and the basic filter (blue). Coarse instructions were provided stochastically 50% of the time. Dynamics noise ranges from 10% to 30%, and observation noise from 10% to 20%. The experiment was run on 5 different test scenarios, each repeated 25 times, then the results were averaged. The result is statistically significant with p-value of 0.1.

operation of this algorithm to real data drawn from tracked vessels in a harbour environment. We show that the proposed algorithm is able to perform much better than a more conventional particle filtering procedure through the use of the hierarchy, and also that it is able to make use of observations that are presented in a form that would be hard to reconcile with the way conventional particle filtering schemes are constructed. We view this work as a step towards systems with more flexible predictive modelling ability in interactive settings, something that is becoming increasingly more prevalent as robots cohabit human-centred environments.

References

1. Isard, M., Blake, A.: Condensation - conditional density propagation for visual tracking. *International Journal of Computer Vision* **29**(1) (1998) 5–28
2. Doucet, A., De Freitas, N., Gordon, N., eds.: *Sequential Monte Carlo methods in practice*. Springer Berlin Heidelberg (2001)
3. Thrun, S., Burgard, W., Fox, D.: *Probabilistic robotics*. MIT press (2005)
4. Tellex, S., Kollar, T., Dickerson, S., Walter, M.R., Banerjee, A.G., Teller, S.J., Roy, N.: Understanding natural language commands for robotic navigation and mobile manipulation. In: *AAAI*. (2011)
5. Kollar, T., Tellex, S., Roy, D., Roy, N.: Grounding verbs of motion in natural language commands to robots. In Khatib, O., Kumar, V., Sukhatme, G., eds.: *Experimental Robotics*. Volume 79 of Springer Tracts in Advanced Robotics. Springer Berlin Heidelberg (2014) 31–47

6. Kuipers, B.: The spatial semantic hierarchy. *Artificial intelligence* **119**(1) (2000) 191–233
7. Belta, C., Bicchi, A., Egerstedt, M., Frazzoli, E., Klavins, E., Pappas, G.: Symbolic planning and control of robot motion [grand challenges of robotics]. *Robotics Automation Magazine, IEEE* **14**(1) (March 2007) 61–70
8. Burridge, R.R., Rizzi, A.A., Koditschek, D.E.: Sequential composition of dynamically dexterous robot behaviors. *The International Journal of Robotics Research* **18**(6) (1999) 534–555
9. Verma, V., Thrun, S., Simmons, R.: Variable resolution particle filter. In: *IJCAI*. (2003) 976–984
10. Brandao, B.C., Wainer, J., Goldenstein, S.K.: Subspace hierarchical particle filter. In: *Computer Graphics and Image Processing, 2006. SIBGRAPI'06. 19th Brazilian Symposium on*, IEEE (2006) 194–204
11. Shabat, G., Shmueli, Y., Bermanis, A., Averbuch, A.: Accelerating particle filter using randomized multiscale and fast multipole type methods. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **PP**(99) (2015) 1–1
12. MacCormick, J., Isard, M.: Partitioned sampling, articulated objects, and interface-quality hand tracking. In Vernon, D., ed.: *Computer Vision ECCV 2000*. Volume 1843 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg (2000) 3–19
13. Yang, C., Duraiswami, R., Davis, L.: Fast multiple object tracking via a hierarchical particle filter. In: *Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on*. Volume 1. (Oct 2005) 212–219 Vol. 1
14. Widynski, N., Mignotte, M.: A multiscale particle filter framework for contour detection. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **36**(10) (Oct 2014) 1922–1935
15. Müllner, D.: Modern hierarchical, agglomerative clustering algorithms. *arXiv preprint arXiv:1109.2378* (2011)
16. Eiter, T., Mannila, H.: Computing discrete fréchet distance. *Rapport technique num. CD-TR* **94** (1994) 64
17. Carlsson, G.: Topology and data. *Bulletin of the American Mathematical Society* **46**(2) (2009) 255–308
18. Pokorný, F.T., Hawasly, M., Ramamoorthy, S.: Multiscale topological trajectory classification with persistent homology. In: *Proceedings of Robotics: Science and Systems*, Berkeley, USA (July 2014)
19. Pokorný, F.T., Hawasly, M., Ramamoorthy, S.: Topological trajectory classification with filtrations of simplicial complexes and persistent homology. *The International Journal of Robotics Research* (2015)
20. Lingala, N., Perkowski, N., Yeong, H., Namachchivaya, N.S., Rapti, Z.: Optimal nudging in particle filters. *Probabilistic Engineering Mechanics* **37**(0) (2014) 160 – 169
21. Lingala, N., Sri Namachchivaya, N., Perkowski, N., Yeong, H.C.: Particle filtering in high-dimensional chaotic systems. *Chaos: An Interdisciplinary Journal of Nonlinear Science* **22**(4) (2012) –
22. Xu, R., Wunsch, D., I.: Survey of clustering algorithms. *Neural Networks, IEEE Transactions on* **16**(3) (May 2005) 645–678
23. Doucet, A., Johansen, A.M.: A tutorial on particle filtering and smoothing: Fifteen years later. *Handbook of Nonlinear Filtering* **12**(656-704) (2009) 3
24. NOAA/BOEM: Marinecadastre.gov. <http://marinecadastre.gov> Accessed: July 2016.